Terms used **WORM write once read many memory pointer** Found **179** of **150,138**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 179 Result page: **1** 2 3 4 5 6 7 8 9 next

Relevance scale

1 Comparison of access methods for time-evolving data

Betty Salzberg, Vassilis J. Tsotras

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Full text available: pdf(529.53 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

Keywords: I/O performance, access methods, structures, temporal databases

2 Process migration

September 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 3

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

3 Implementing multidestination worms in switch-based parallel systems: architectural alternatives and their impact

Craig B. Stunkel, Rajeev Sivaram, Dhabaleswar K. Panda

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2

Additional Information: full citation, abstract, references, citings, index

Full text available: pdf(1.98 MB) terms

Multidestination message passing has been proposed as an attractive mechanism for efficiently implementing multicast and other collective operations on direct networks. However, applying this mechanism to switch-based parallel systems is non-trivial. In this paper we propose alternative switch architectures with differing buffer organizations to implement multidestination worms on switch-based parallel systems. First, we discuss issues related to such implementation (deadlock-freedom, replicatio ...

4 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

5 A taxonomy of computer program security flaws

Carl E. Landwehr, Alan R. Bull, John P. McDermott, William S. Choi

September 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 3

Full text available: pdf(3.81 MB) Additional Information: full citation, abstract, references, citings, index terms, review

An organized record of actual flaws can be useful to computer system designers, programmers, analysts, administrators, and users. This survey provides a taxonomy for computer program security flaws, with an Appendix that documents 50 actual security flaws. These flaws have all been described previously in the open literature, but in widely separated places. For those new to the field of computer security, they provide a good introduction to the characteristics of security flaws and how they ...

Keywords: error/defect classification, security flaw, taxonomy

6 A pipelined memory architecture for high throughput network processors

Timothy Sherwood, George Varghese, Brad Calder

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available: pdf(213.66 KB) Additional Information: full citation, abstract, references, citings

Designing ASICs for each new generation of backbone routers is a time intensive and fiscally draining process. In this paper we focus on the design of a programmable architecture for backbone routers, based on the manipulation of wide irregular memory words, that can provide a feasible design alternative to custom ASICs. We propose a pipelined memory design that emphasizes worst-case throughput over latency, and co-explore architectural tradeoffs with the design of several important network algo ...

7 File organizations and access methods for CLV disks

S. Christodoulakis, D. A. Ford

May 1989 **ACM SIGIR Forum , Proceedings of the 12th annual international ACM SIGIR conference on Research and development in information retrieval**, Volume 23 Issue 1-2

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

A large and important class of optical disc technology are CLV format discs such as CD ROM and WORM. In this paper, we examine the issues related to the implementation and performance of several different file organizations on CLV format optical discs such as CD ROM and WORM. The organizations examined are based on hashing and trees. The CLV

recording scheme is shown to be a good environment for efficiently implementing hashing. Single seek access and storage utilization levels a ...

8 eNVy: a non-volatile, main memory storage system

Michael Wu, Willy Zwaenepoel

November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29 , 28 Issue 11 , 5

Full text available: pdf(1.32 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture of eNVy, a large non-volatile main memory storage system built primarily with Flash memory. eNVy presents its storage space as a linear, memory mapped array rather than as an emulated disk in order to provide an efficient and easy to use software interface.Flash memories provide persistent storage with solid-state memory access times at a lower cost than other solid-state technologies. However, they have a number of drawbacks. Flash chips are ...

9 Demonic memory for process histories

P. R. Wilson, T. G. Moher

June 1989 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1989 Conference on Programming language design and implementation**, Volume 24 Issue 7

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

Demonic memory is a form of reconstructive memory for process histories. As a process executes, its states are regularly checkpointed, generating a history of the process at low time resolution. Following the initial generation, any prior state of the process can be reconstructed by starting from a checkpointed state and re-executing the process up through the desired state, thereby exploiting the redundancy between the states of a process and the description of that process (i.e., a comput ...

10 A high performance multi-structured file system design

Keith Muller, Joseph Pasquale

September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles**, Volume 25 Issue 5

Full text available: pdf(1.40 MB) Additional Information: full citation, abstract, references, citings, index terms

File system I/O is increasingly becoming a performance bottleneck in large distributed computer systems. This is due to the increased file I/O demands of new applications, the inability of any single storage structure to respond to these demands, and the slow decline of, disk access times (latency and seek) relative to the rapid increase in CPU speeds, memory size, and network bandwidth.We present a *multi-structured file system* designed for high bandwidth I/O and fast response. Our design ...

11 Interactive Editing Systems: Part II

Norman Meyrowitz, Andries van Dam

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: pdf(9.17 MB) Additional Information: full citation, references, citings, index terms

12 Generative communication in Linda

David Gelernter

January 1985 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 7 Issue 1

Full text available: pdf(2.48 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Generative communication is the basis of a new distributed programming langauge that is intended for systems programming in distributed settings generally and on integrated network computers in particular. It differs from previous interprocess communication models in specifying that messages be added in tuple-structured form to the computation environment, where they exist as named, independent entities until some process chooses to receive them. Generative communication results in a number ...

13 Methods for message routing in parallel machines

Tom Leighton

July 1992 **Proceedings of the twenty-fourth annual ACM symposium on Theory of computing**

Full text available: pdf(2.32 MB) Additional Information: full citation, references, citings, index terms

14 The LHAM log-structured history data access method

Peter Muth, Patrick O'Neil, Achim Pick, Gerhard Weikum

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Full text available: pdf(494.76 KB) Additional Information: full citation, abstract, index terms

Numerous applications such as stock market or medical information systems require that both historical and current data be logically integrated into a temporal database. The underlying access method must support different forms of "time-travel" queries, the migration of old record versions onto inexpensive archive media, and high insertion and update rates. This paper presents an access method for transaction-time temporal data, called the log-structured history data access method (L ...

Keywords: Data warehouses, Index structures, Performance, Storage systems, Temporal databases

15 Hardware/software tradeoffs in a variable word width, variable queue length buffer memory

A. C. Parker, A. W. Nagle

March 1977 **ACM SIGARCH Computer Architecture News , Proceedings of the 4th annual symposium on Computer architecture**, Volume 5 Issue 7

Full text available: pdf(401.04 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes a buffer memory design which is integratable on a single 40 pin package. The memory consists of four variable word width, variable length queues, address logic, and logic for configuring the data widths and queue lengths. Simple program commands input on the data lines reconfigure the memories when the circuit is in program mode. In data mode, WRITE and READ requests to a particular buffer enable the memory to load and access itself, update its pointers, and check for f ...

16 Anatomy of a message in the Alewife multiprocessor

John Kubiatowicz, Anant Agarwal

August 1993 **Proceedings of the 7th international conference on Supercomputing**

Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, references, citings, index terms

Shared-memory provides a uniform and attractive mechanism for communication. For

efficiency, it is often implemented with a layer of interpretive hardware on top of a message-passing communications network. This interpretive layer is responsible for data location, data movement, and cache coherence. It uses patterns of communication that benefit common programming styles, but which are only heuristics. This suggests that certain styles of communication may benefit from direct access to the ...

17 A new switch chip for IBM RS/6000 SP systems

Craig B. Stunkel, Jay Herring, Bulent Abali, Rajeev Sivaram

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(177.66 KB) Additional Information: full citation, references, citings, index terms

18 Applications II: Time and area efficient pattern matching on FPGAs

Zachary K. Baker, Viktor K. Prasanna

February 2004 **Proceeding of the 2004 ACM/SIGDA 12th international symposium on Field programmable gate arrays**

Full text available: pdf(354.50 KB) Additional Information: full citation, abstract, references, index terms

Pattern matching for network security and intrusion detection demands exceptionally high performance. Much work has been done in this field, and yet there is still significant room for improvement in efficiency, flexibility, and throughput. We develop a novel linear-array string matching architecture using a buffered, two-comparator variation on the Knuth-Morris-Pratt(KMP) algorithm. For small (16 or fewer characters) patterns, it competes favorably with the state-of-the-art while providing bett ...

Keywords: Knuth-Morris-Pratt, network intrusion detection, string matching

19 Deciding when to forget in the Elephant file system

Douglas S. Santry, Michael J. Feeley, Norman C. Hutchinson, Alistair C. Veitch, Ross W. Carton, Jacob Ofir

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles**, Volume 33 Issue 5

Full text available: pdf(1.61 MB) Additional Information: full citation, abstract, references, citings, index terms

Modern file systems associate the deletion of a file with the immediate release of storage, and file writes with the irrevocable change of file contents. We argue that this behavior is a relic of the past, when disk storage was a scarce resource. Today, large cheap disks make it possible for the file system to protect valuable data from accidental delete or overwrite.This paper describes the design, implementation, and performance of the Elephant file system, which automatically retains all impo ...

20 COMA: an opportunity for building fault-tolerant scalable shared memory multiprocessors

Christine Morin, Alain Gefflaut, Michel Banâtre, Anne-Marie Kermarrec

May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf(1.30 MB) Additional Information: full citation, abstract, references, citings, index terms

Due to the increasing number of their components, Scalable Shared Memory Multiprocessors (SSMMs) have a very high probability of experiencing failures. Tolerating node failures therefore becomes very important for these architectures particularly if they must be used

for long-running computations. In this paper, we show that the class of Cache Only Memory Architectures (COMA) are good candidates for building fault-tolerant SSMMs. A backward error recovery strategy can be implemented without sign ...

Keywords: Scalable Shared

Results 1 - 20 of 179 Result page: **1** 2 3 4 5 6 7 8 9 next